An Agent-Based Approach to Aircraft Conflict Resolution with Constraints

Karen A. Harper,[*] Sean L. Guarino,[†] Aaron White,[‡] and Mark L. Hanson[§]
Charles River Analytics
Cambridge, MA 02138

Karl D. Bilimoria[**] and Daniel Mulfinger[††]
NASA Ames Research Center
Moffett Field, CA 94035

ABSTRACT

In recent years, much attention has been given to the development of emerging concepts of decentralized control within the National Airspace System (NAS). The introduction of a distributed decision-making environment in the ATM system will allow individual users (pilots and airline operations centers (AOC)) more freedom in changing and optimizing flight plans while enroute in comparison to the current pre-flight planning that happens now. Such an approach to ATM system operations will require individual decision-makers within the airspace to identify and solve routing problems due to traffic conflicts, weather avoidance, and Special Use Airspace (SUA) dynamics, in real time. In this paper, we apply a modeling and simulation based approach to investigate the potential for a *principled negotiation*-based approach to the distributed ATM problem. The primary objective is to identify the frequency with which truly collaborative (i.e., negotiation) behavior may be required within realistic ATM operations.

---

[*] Principal Scientist; kharper@cra.com. Member, AIAA
Author to whom all correspondence should be addressed
[†] Software Engineer. Member, AIAA
[‡] Software Engineer
[§] Senior Scientist. Member, AIAA
[**] Research Scientist, Automation Concepts Research Branch. Associate Fellow, AIAA
[††] Software Engineer, Raytheon ITSS

INTRODUCTION

Emerging concepts for future air traffic management (ATM) systems and procedures will dramatically change human roles and tasks in the National Airspace System (NAS). The introduction of a distributed decision-making environment in the ATM system will allow individual users (pilots and airline operations centers (AOC)) more freedom in changing and optimizing flight plans while enroute in comparison to the current pre-flight planning that happens now (FAA, 1995). By giving pilots and AOCs the capability to reroute while enroute, operations could be considerably streamlined, and significant improvements could be realized in fuel efficiency and schedule maintenance. The introduction of this envisioned distributed decision-making (DDM) environment, however, will have profound implications on pilot/controller/AOC information requirements, roles and responsibilities, allocation of workload, communication, and decision-making *throughout* the ATM system.

NASA and the FAA are focusing technical initiatives in the development of new systems-level concepts for ATM operations, to meet the projected airspace demands of the future, while maintaining the overall safety of the NAS. One such concept, *free flight* (RTCA, 1995) supported total autonomy on the part of pilots within the airspace, giving them full authority to self-separate, and perform enroute flight planning and replanning independent of air traffic control. An evolved version of free flight, known as *distributed air/ground traffic management (DAG-TM)* (NASA, 1999) reconsidered the role of air traffic control in the future airspace and developed a concept of ATM operations that shares the decision-making authority among pilots *and* controllers, who are expected to behave cooperatively to optimize both individual and global operations, while maintaining safety as the highest priority. Further efforts to investigate potential operational concepts are currently being pursued by NASA (Denery & Weathers, 2001), through a research path addressing three primary objectives:

1) To develop the tools and modeling capabilities for assessing the requirements of an advanced air transportation system;
2) To conduct system-level assessments of new capabilities for ATM; and
3) To develop the core technologies required to complete the goals of the Free Flight Program initiatives and set the foundations for the air transportation system beyond the Free Flight Program as currently defined.

The ultimate objective of these ATM system redesign efforts is to generate new rules and protocols governing individual behavior of the key decision-makers in the air transportation system, including pilots, air traffic controllers and airline dispatchers. Because these decision-makers will ultimately drive overall ATM system performance and safety, any M&S approach to system analysis *must* include realistic human behavior representations (HBRs) of the key decision-makers. Furthermore, these HBRs must incorporate the underlying perceptual and cognitive capabilities and limitations that determine decisions made by these key players, if we are to capture realistic behaviors in the DDM environment.

If computational agent-based instantiations of these HBRs can be successfully embedded in an appropriate ATM simulation, we will then be in a position to study the effects that specific DDM behaviors have on the overall performance of the ATM system. This calls for significant flexibility in the HBR representations themselves, so that we can adjust individual behaviors and observe the effects that such adjustments have on overall system performance. With this capability, we will then be able to identify the combined set of individual behaviors, procedures and protocols that optimize ATM system-level performance, with respect to safety and efficiency. In effect, we will have an M&S tool that will allow us to evaluate emergent ATM performance/safety as a function of individual decisions made by the players.

To support these objectives, this work has been focused on the development of an agent-based modeling approach of human behavior, which allows us to investigate the feasibility of distributed control through *principled negotiation* (Jacolin & Stengel, 1998; Wangermann & Stengel, 1996) in the air traffic environment, and provides a testbed for new communications protocols, information requirements, and traffic management models. The principled negotiation approach to distributed decision-making involves the search for conflict resolution maneuvers that provide the greatest mutual gains with respect to each participating individual's requirements. Following from previous efforts (Harper, Mulgund, Zacharias & Kuchar, 1998; Harper, Mulgund, Guarino, Mehta & Zacharias, 1999), we have furthered the development of agent-based models of pilot, air traffic controller, and airline dispatcher behavior. The resulting agents are responsible for collaborating in a simulated advanced ATM environment to resolve traffic conflicts, weather problems, and restricted airspace avoidance issues in a safe, fair and timely manner.

Under previous development efforts, we have demonstrated the capability of the agent-based models of collaborative decision-making to identify specific problems and solve them through democratic negotiation. In this paper, we are interested in now investigating the actual requirements for truly collaborative decision-making in the ATM environment. Algorithmic conflict resolution and rerouting methods have been developed and demonstrated in the literature, and could be implemented within a real-time automated route planning and adjustment tool (Bilimoria, 2000; Krozel, Mueller & Hunter, 1996; Tomlin & Pappas, 1998). The question now becomes whether these automated methods are sufficient to support separation maintenance and hazard avoidance, or if more involved distributed control strategies are required to maintain the efficiency and safety of the NAS.

The paper is outlined as follows. First, we provide the technical background underlying the development of our agent-based models of human decision-making and collaboration through principled negotiation. We then present an overview of the knowledge engineering (KE) based model development process employed in the construction of agent-based representations of commercial pilots, air traffic controllers, and AOC dispatchers, followed by a summary of the results of the KE efforts. This discussion concludes in a description of the information processing, situation assessment, and collaborative decision-making functions developed across the three targeted agent models. We then present an overview of the integrated system architecture that brings these models together within the simulated air traffic environment provided by the Future ATM Concepts Evaluation Tool (FACET) (Bilimoria, Srindhar, Chatterji, Sheth & Grabbe, 2000). To support the experimental objective of this paper, we then describe a simulated air traffic scenario that incorporates realistic traffic, weather, and SUA dynamics for Jacksonville Center (ZJX). We populate this scenario with dynamic agent representations and analyze their collaborative problem-solving behavior to determine the level of complexity that can be expected in a distributed ATM environment, and the frequency and nature with which enroute flight adjustments must be truly negotiated among multiple decision-makers within the system. Finally, we present some conclusions and recommendations for further study.

BACKGROUND

Developing agent models of collaborative decision-making behavior among commercial pilots, air traffic controllers and AOC dispatchers are based on SAMPLE (Situation Assessment Model of Pilot-in-the-Loop Evaluation), an agent-based human behavior modeling architecture originally developed to model tactical aviation pilots, but designed within a domain-independent architecture. SAMPLE is based on a hierarchical definition of behavior, where the highest levels represent ongoing cognitive tasks and the lowest levels define specific information processing, situation assessment, and procedurally-driven decision-making models. We believe that SAMPLE provides a strong agent-based modeling approach for advanced ATM environments, supporting the modular development of various agent representations.

SAMPLE Overview

The SAMPLE human behavior model is based on the Rasmussen Hierarchy of human information processing and skilled behavior (Rasmussen, 1983; Rasmussen, 1986), which provides a strong unifying theoretical framework for analysis of different human skills that may be modeled within a real-time simulation. By dividing skilled behavior into categories based on the degree of automaticity, complexity, and level of cognitive processing, this framework supports systematic skill decomposition and measurement of individual aspects of the overall skill on a part-task basis.

SAMPLE implements the Rasmussen analysis in a system that has its roots in the Optimal Control Model (OCM) (Kleinman, Baron & Levison, 1970). This is an information-processing model of the operator of a dynamic system, grounded in modern control and estimation theory, which accounts for closed-loop human/machine performance across a range of primarily continuous control tasks (e.g., flight-path control). The first attempt to integrate continuous and discrete information processing with the structured procedural activities of the flight crew was with PROCRU, a model developed to evaluate commercial approach procedures during landing (Baron, Muralidharan, Lancraft & Zacharias, 1980; Milgram, van der Wijngaart, Veerbeek, Bleeker & Fokker, 1984). Here, the task of situation assessment was first made explicit, and linked to procedurally-driven rules of engagement. Finally, an enhanced model known as the crew/system integration model (CSIM) has been used for the analysis of anti-aircraft artillery crews (Zacharias, Baron, Muralidharan & Kastner, 1981), fighter attack missions (Zacharias & Baron, 1982), and in a precursor effort to the original SAMPLE development effort (Zacharias, Miao, Illgen, Yara & Siouris, 1996).

The SAMPLE architecture combines elements of the Rasmussen Hierarchy and the CSIM model. Figure 1 decomposes the SAMPLE model into its constituent conceptual components:

1) An **Information Processing** module that processes information generated by the simulation, to yield system states and event cues;

2) A **Situation Assessment** module that uses event cues to generate the current assessed situation;
3) A **Decision Making and Procedure Execution** module that selects among alternative procedures to produce control actions, based on the current situation and estimated states.

The front end to the agent model's SA-based decision-making model is the Information Processing system, which transforms fused sensor data into situationally relevant semantic variables. There are three algorithmic subcomponents to this system: Attention Allocation, Sensory Processing, and Perceptual Processing.

The **Attention Allocation** subcomponent accounts for the human's sensory limitations and defines attention allocation among competing sources of information. This subcomponent reflects the human's inability to process all sources of information simultaneously, and must, therefore, decide on which sources to attend to. Appropriate allocation of attention is considered a function of the current tasks defined by the Decision-Making function downstream, so that more attention is allocated to the information resources that are directly relevant to the current tasks and procedures.

The **Sensory Processing** subcomponent represents the physical sensor modeling in SAMPLE. This data includes an overall view of the world, including accumulated data from visual, audio, and, for some agents, a host of other sources. In the flight domain, this might include various sensor readings supplying information about weather or traffic conditions, along with communications received from ATC. This data is filtered according to information requirements specified by attention allocation.

The **Perceptual Processing** subcomponent acts as a domain-relevant interpreter for the sensory information output, searching the available information for specific events relevant to the environment and current situation, and fusing this data into a host of relative perceived events. For example, one might search all incoming messages to determine if an expected message has been received, so that message can be interpreted as needed.

The **Situation Assessment** module generates a high-level interpretation of the operational situation, as a function of perceived events and situational memory. Situation assessment is performed via a cognitive matching of low-level detected events and high-level representations of the current state in a context-sensitive manner. Examples of tasks performed by this system might include air traffic conflict detection, threat identification and prioritization.

The assessed situation feeds the **Decision Making and Procedure Execution** module. The **Decision Making** subcomponent emulates a human's rule-based decision-making behavior to select a procedure to implement, feeding rules available in the decision rule memory with the overall assessed situation. The resulting decisions feed the attention allocator, defining the information required to support the selected procedure implementations.

The **Procedure Execution** subcomponent translates high-level decisions into the required low-level motor actions or communications that implement selected procedures. Details of the actions that make up a high-level procedure are maintained in the procedure memory. One important task performed in procedure execution in any multi-agent system is the inter-agent communications. The SAMPLE inter-agent communication model is based on the military Command and Control Simulation Interface Language (CCSIL), as this was directly relevant in the original domain for which SAMPLE was developed, modeling air-to-air combat pilots. In the future, a more generic approach to inter-agent communications (e.g. the Knowledge Query Manipulation Language, or KQML (Finn, Labrou & Peng, 1998)) may be integrated into the architecture.

Agent-Based Collaborative Decision-Making via Principled Negotiation

Though several approaches to the collaborative decision-making process could prove useful in advanced ATM modeling (Guarino, Harper, White, Omartian & Zacharias, 2001), we selected a particular form of negotiation through bargaining to define the ATM system configuration in a decentralized control paradigm. *Principled negotiation* (Wangermann & Stengel, 1996) considers the individual needs and preferences of all agents involved in a decision-making task. It relies on the notion of *mutual gain* in defining options: agents drive their search for acceptable solutions by considering not only their own needs and preferences (e.g., aircraft performance constraints, or minimizing delays and fuel burn), but also those of other agents affected by the solution and overall system constraints (e.g., restricted airspace). By proposing solutions and reporting acceptances or rejections, agents bargain to reach an acceptable solution. For example, consider the case where a pair of aircraft must perform a strategic conflict avoidance maneuver. The principled negotiation model would have both pilots negotiating a solution that would not only avoid a protected zone conflict, but would also minimize each aircraft's deviation from its original flight plan.

Principled negotiation is carried out through an iterative optimization method within each agent's decision-making model, as shown in Figure 2 (Wangermann & Stengel, 1996). One of four outcomes are possible:

1) **Propose an option**: An agent searches for a solution beneficial to itself and poses it to the group;
2) **Accept the option**: All other agents review the proposed option and accept it if it is considered better than any currently considered option;
3) **Reject the option**: The proposed solution is rejected if it is considered less beneficial that another option by any agent;
4) **Instruct agents to implement the option**: Once an option has been accepted by all agents, the proposing agent instructs the group to proceed with its implementation

Jacolin and Stengel (Jacolin & Stengel, 1998) present a division of negotiating behavior between *maximizing agents* and *satisficing agents*. Maximizing agents search for solutions that are not only acceptable but are optimal with respect to some utility function. Satisficing agents, however, do not try to optimize in any sense. Rather, they accept any presented solution that satisfies some constraint boundaries, but do not optimize performance within those boundaries.

Jacolin and Stengel (1998) suggest that, in the ATM domain, pilots and airlines would likely act as maximizing agents, attempting to optimize performance according to cost in time, fuel, etc. ATC, however, would likely behave as a satisficing agent, accepting any proposed solution that meets the requirements of overall system safety. We implemented the Wangermann and Stengel model of principled negotiation within our agent-based representations of pilot, ATC, and airline dispatcher behavior, and extended the model to adjust the negotiating behavior of the agents as a function of ATM system configuration parameters. These extensions to the model are intended to represent the heterogeneous nature of individuals within the ATM system, in terms of roles and responsibilities, as well as capability (as a function of equipage, for example).

MODELING APPROACH

We now present the results of our agent-based modeling approach for distributed decision-making within the ATM system of the future. We first describe our approach to model development through knowledge engineering (KE), and then present the results of KE exercises to represent the distributed roles and responsibilities of the key players in the ATM system.

Model Development via Knowledge Engineering

Developing computational cognitive models via knowledge engineering is a two-stage process. The first stage consists of a Cognitive Task Analysis (CTA), which takes a problem or domain area and extracts the relevant SA information requirements. These requirements are defined as the dynamic information needs associated with the major goals and sub-goals that pilots, ATC, and airline dispatchers must perform. To identify these requirements, we applied a goal-directed task analysis based on the methodology of Endsley (1993). This approach involves four steps: 1) identify the key decision makers associated with the domain (e.g. pilot, ATC, and dispatcher); 2) identify the major goals for each decision-maker, along with the major sub-goals necessary for meeting each of these goals; 3) identify the primary decisions that needed to be made for each sub-goal; and 4) identify the SA information requirements (e.g., data needed, how information is integrated, combined, etc…) for making these decisions and carrying out each sub-goal.

The second stage of model development uses the identified SA information requirements and integrates it within the simulation environment, accounting for data availability (e.g., FACET provides access to flight plan information but no knowledge of runway conditions at destination airports) to construct a Computational Cognitive Model. Effectively, this effort is geared towards translating the CTA results into a computationally accessible representation. In our case, the goal is to generate SAMPLE component-based representations of the information processing, situation assessment, and decision-making tasks dictated by the CTA results. There are three aspects to this translation. First, the information requirements must be analyzed to determine the key events and associated information processing needs. For example, raw distance measures must be converted into semantically relevant measures such as "below minimums," "near minimums," and "above minimums", which humans typically use.

Next, computational SA models are constructed to process the events to generate situation awareness. This involves analyzing the assessment and determining what is the most appropriate technology or algorithm (e.g. Bayesian belief networks) to relate the assessment (e.g. weather risk) to the identified events (e.g. thunderstorm type and distance to flight path). Considerations include robustness, uncertainty levels, and available processing capability.

The final module performs decision-making based on the assessed situation. These decisions are the same ones identified in the CTA, so the objective here is to translate those decisions into computational representations of procedure implementations.

It should be noted that validation and verification of the CTA results, and their corresponding transition into computational components within the SAMPLE architecture should be pursued through follow-up experiments with a population of SMEs. While this effort has not included a significant validation of computational processes within the agent-based models, we recommend that future work include such a component.

Knowledge Engineering Results

The two-stage CTA and Computational Cognitive Model development processes were applied to the advanced ATM problem. First, the key decision makers within the advanced ATM operational environment were identified. They are: 1) airline pilots; 2) air traffic controllers; and 3) airline operations center (AOC) dispatchers. Together, these three entities form a triad whose overriding goal is to safely and efficiently operate within the NAS. To identify the major goals and sub-goals along with the major decisions and SA requirements for each group, knowledge elicitation interviews were conducted with practicing members of each group. In the case of pilots and ATCs, existing SA analysis by Endsley (1998; 1994) was used and supplemented with additional interviews and visits to actual facilities (for ATC). We, therefore, focused on applying the aforementioned two-stage model development via a knowledge engineering process to identify the information needs of AOC dispatchers. The final results comprise three computational cognitive models, which model human decision-making behavior for pilots, ATCs, and AOC dispatchers. We now present summary of results from Endsley for pilots (1998) and ATCs (1994) along with a more detailed description for AOC dispatchers that directly resulted form this effort.

KE Results for Commercial Airline Pilots

Commercial airline pilots are in command of the aircraft and crew, and are responsible for the safety of the passengers and plane. Federal Aviation Regulation (FAR) section 121.533 states that "during flight, pilots have full control and authority in the operation of the aircraft without limitation." Based on many hours of interviews, Endsley (1998) found that pilots have one overriding goal of safely transporting the plane from the origin to the destination. As shown in Figure 3, Endsley found that pilots have four major sub-goals, namely 1) select best path to destination; 2) execute desired flight path safely, efficiently, and with acceptable ride comfort; 3) manage resources effectively; and 4) satisfy the customer.

To study agent-based models of air traffic management, we focused on addressing the CTA-identified major decisions and SA requirements associated with the first two sub-goals. With respect to selecting the best path to the destination, our agents make decisions and require SA information regarding assessing a flight plan ("is the flight plan safe?"), determining changes to the flight plan ("how should I maneuver to avoid a traffic conflict?"), and re-planning the flight path. With respect to executing a desired flight plan, the agents evaluate and execute a flight plan, avoid conflicts ("am I in conflict with another plane?"), avoid hazardous weather ("am I in conflict with weather?"), and minimize the impacts of abnormal ATC situations ("fly around a special use airspace"). A pilot agent also addresses requirements associated with the third sub-goal shown in Figure 3. For example, a pilot agent may inform its corresponding AOC dispatcher of any potential flight plan changes if they are considered significant and the pilot deems it necessary to consult with the dispatcher.

KE Results for ATC

Air traffic controllers are responsible for managing traffic flows and keeping airplanes within established minimum separation levels. Similar to commercial airline pilots, Endsley (1994) found that ATCs have one overriding goal – assure flight safety. Figure 4 shows the corresponding sub-goals, which ATCs try to achieve: 1) avoid conflicts; 2) provide flight service; and 3) handle perturbations.

From a CTA perspective, we focused primarily on avoiding conflicts and to a lesser degree providing flight service and handling perturbations. Thus, we include some of the major decisions and SA information requirements associated with separating airplanes ("are the vertical or lateral separation requirements met?" and "what is the time to separation loss?") and avoiding airspace conflicts ("what airspace should be avoided and which pilots need to be contacted?"). To address the issue of providing flight service and handling perturbations, we focused on implementing decision-making capability to determine the impact that a potential maneuver will have on other entities within the airspace and when a sector should be locked due to traffic density.

KE Results for Airline Operations Center (AOC) Dispatchers

An AOC dispatcher is the third player in the FAR specified concept of the operations triad. Whereas ATC are specifically tasked to avoid traffic conflicts and pilots are tasked to operate the vehicle from the origin to the destination, the dispatcher is charged with developing a safe flight plan and monitoring all risks and hazards that may jeopardize or compromise the safety of the flight plan.

Similar to ATCs and pilots, a dispatcher's major goals and sub-goals are directly related to the FARs. In particular, the identified major goal of the dispatcher is the creation and maintenance of a safe plan of operation for each flight, as dictated by FAR 121.533 - Responsibility for Operation Control: Domestic Operations. Based on the FARs, our KE study produced a cognitive task model that consists of three major sub-goals that contribute to the single major goal of creating and maintaining a safe plan of operation. As shown in Figure 5, the three major sub-goals are: 1) initiate a plan; 2) conduct a flight; and 3) terminate a flight.

With respect to *Initiating a Plan*, the dispatcher first creates a plan that avoids hazards, is cost-effective, and meets federal rules and regulations. Then, according to FAR 121.533, the dispatcher and pilot must jointly agree on the plan. In creating the plan, the dispatcher uses metrological data to forecast thunderstorms (TS), clear air turbulence (CAT), low-ceiling visibility, low-level adverse wind conditions, and icing conditions to develop a safe route. These predictions are used to answer key questions such as "should the plane depart?", "is the chosen route safe and acceptable?", "is the landing site acceptable?", and so on. The dispatcher must also ensure that the chosen plan meets federal rules and regulations regarding such items as takeoff and landing alternates, e.g. "is a destination alternate required?", as well as requirements regarding minimum fuel, acceptable takeoff weight, and performance and certification limitations.

Once the aircraft is airborne, the dispatcher's primary sub-goal is *Conducting the Flight*. This is accomplished via two primary tasks, which are monitoring and re-planning the flight. In the monitoring task, the assessments and decisions are made to: 1) assess the pilot's intent ("is the pilot following the flight plan?"); 2) assess the overall system safety ("is the system safe and what is the reason for risk?"); 3) assess the viability of the landing options ("are the runways acceptable for landing?"); 4) handle pilot and passenger requests; 5) evaluate potential ATC re-routes ("what are the results of ATC delays and holding patterns on the flight?"); 6) evaluate the impact of mechanical failures ("can plane safely and legally continue on current flight plan with a given failure?"); and 7) inform the pilot of potential hazardous or risky flight conditions. In the re-planning task, the dispatcher uses the assessments from the monitoring tasks to: 1) determine suitable airport to land (if necessary) and 2) generate a plan to eliminate or minimize the effects of the identified hazards, risks, and system changes on the flight.

The third and final sub-goal is *Flight Termination*. Here, the dispatcher analyzes risk and hazard data to determine if the flight should be terminated due to unsafe conditions, which may include weather or instrument or equipment failure. According to FAR 121.627, the dispatcher may not allow a flight to continue if in the opinion of the dispatcher the flight cannot be operated safely. As indicated in Figure 5, our KE analysis did not focus on flight termination.

Software System Architecture

Figure 6 presents the overall architecture of our proposed agent-based model of pilot, ATC, and airline dispatcher performance under advanced ATM operations, building directly on our previous prototype design (Harper et al., 1999). The system contains the following key components:

- Analytical models of the pilot, air traffic controller (ATC), and airline dispatcher agents, representing the key activities of air traffic situation assessment, collaborative decision-making and plan execution
- The ATM executive, which models the NAS environment (airspace configuration, flight rules, separation criteria, etc.)
- The FACET air traffic simulator, which models the dynamics of each represented aircraft

The pilot agent is instantiated *N* times, to represent *N* aircraft operating within the simulation. Each of these aircraft communicates with other aircraft and with one ore more ATC and dispatcher agents via simulated voice and datalink, as required. The problem of multi-agent air traffic conflict resolution is handled by a simulated *negotiation process* among agents affected by a potential conflict or weather hazard.

The agent architecture contains three key components, as follows:

1) **Air Traffic Situation Assessor**: In high-stress, time-critical situations, expert decision-makers focus their attention on the task of assessing the situation based on information gathered from available resources (Klein, 1989). In the ATM domain, this entails the collection of information about potential traffic hazards from sensor data and known or

requested intent information, and the fusion of that information into a global picture of the current and projected state of the air traffic sector (either in the vicinity of an aircraft or in some fixed region of airspace).

2) **Collaborative Decision-Making Module**: Once individual agents have developed an understanding of the current and projected air traffic situation, any predicted conflict problems must be solved. Since the air traffic system is highly coupled (i.e., actions by one aircraft affect many others), the conflict resolution process must be carried out in a distributed manner, such that all users' interests are represented in the solution, and all users are aware of how the solution will affect their own plans. This is accomplished by carrying out a collaborative decision-making process where everyone's interests are represented (to the degree possible) in the chosen solution. This collaboration is based on the exchange of information throughout the problem-solving exercise by voice and datalink.
3) **Plan Execution**: The plan execution function takes in a flight plan representation and defines the control variables required to adhere to it. For pilots, this entails updating waypoint definitions to include new details of their flight plan. ATC must assess the negotiated solution to identify any aircraft that are affected by it that may not have received the updated plan, and inform them of the updates.

Simulation of Advanced ATM Environment

We now present our approach to integrating the target agent-based models of pilot, air traffic controller, and airline dispatcher behavior within the simulated NAS system provided by FACET. First, we highlight some of the modeling capabilities provided by FACET that were particularly useful in supporting our objectives. We then discuss the added modeling capability that we provided within the FACET simulation environment to model weather cells and restricted airspace, which allowed us to examine our developing agents' capabilities to detect and resolve a wider array of airspace issues. Finally, we discuss our implementation of a high-level "agent manager" entity that support the real-time performance of the integrated system.

Modeling the National Airspace System

The Future ATM Concepts Evaluation Tool (FACET) is an air traffic management research tool developed at NASA Ames Research Center (Bilimoria et al., 2000). It was designed to provide a simulation test environment for the evaluation of advanced ATM concepts, since the architecture supports straightforward integration of air-traffic analysis applications.

As a tool developed for testing advanced ATM concepts, FACET, shown in Figure 7, is an ideal simulation environment for the evaluation of our principled negotiation model for future decentralized ATM operations. FACET incorporates many of the important features of our previous in-house graphical user interface (Guarino, Harper & Zacharias, 2000; Harper et al., 1999). In addition, FACET has some advantages over our in-house system. For instance, it operates on a variety of platforms, including SGI IRIX and Microsoft Windows environments. FACET has a more accurate depiction of the airspace over the continental US, allowing us to implement scenarios of much larger scope, and account for the configuration of the NAS system. Furthermore, it implements low-cost models of aircraft dynamics that allow for large scenarios involving many aircraft capable of real time performance.

The FACET software environment also provided significant flexibility with respect to integrating our agent-based models of human behavior within the ATM system. The existing separation of graphical functionality implemented in Java, and computational functions in C, with a Java Native Interface (JNI) bridge between the two, provided an easily extended software architecture for the integration of our agent models. The SAMPLE agent framework was implemented within a stand-alone C++ library, and a second JNI interface was developed to integrate it within the FACET architecture.

Modeling Restricted Airspace within FACET

One of the primary objectives of the agent development effort was to implement the capability to negotiate not only air traffic conflict resolutions, but also weather avoidance maneuvers. To support this objective, we developed an independent weather representation within the FACET environment.

In choosing a weather representation, we explored many options, trying to maximize realism while still maintaining maximal customizability and computational simplicity. With the ability to adapt the FACET polygon avoidance algorithm, weather represented as polygonal prisms received early and serious consideration. Fortunately, a polygonal prism representation of weather also affords much customization; the number of vertices is variable, simply changing the location of one or more vertices can modify the shape and position, and cell tops and bottoms naturally fit well with the ends of the prism. Polygonal prisms also allow for much simpler calculations than ellipsoids or spherical shapes, thus increasing the real-time performance of the simulated weather representation.

Restricted areas, prohibited areas, military operations areas, and other special use airspaces usually have polygonal boundaries. Given that special use airspace (SUAs), much like severe weather, present dynamic obstacles to safe flight, a common representation for the two would allow for additional modeling of dynamic SUAs within our principled negotiation framework. For these reasons, we elected to use polygonal prisms as a generic "spatial region" that can be modified to represent thunderstorms, clear air turbulence, or SUAs.

In addition to vertical and lateral boundaries, weather and turbulence areas have an associated severity level. For thunderstorms, severity is listed on a ten point Video Integrated Processor (VIP) scale, with severities less than 3.0 appearing as green or light intensity, severities between 3.0 and 5.0 appearing as yellow or moderate intensity, and 5.0 or greater appearing as red or severe intensity. Turbulence regions also use a 10-point scale with the same categorization as thunderstorms, but not a VIP scale since turbulence assessments are based on pilot reports (PIREPs) as opposed to sensor reports. Special use airspace does not require such a severity level, as we assume that all special use airspaces are no-fly zones.

Spatial regions are defined by inputting coordinates, severities, and altitudes in a data file, and then loading it from FACET by selecting CRA Weather from the CRA Applications submenu of the Applications menu, as described in Appendix … The user defines the weather cell's vertices at user-defined simulation times. Therefore, one can easily change the shape and location of a weather cell simply by setting the vertices to new values at a different simulation time in the input file.

Supporting Real-Time Performance

The computational requirements of the developed agent-based representations of human decision-making behavior are intensive. Since the developing agent-based models of human decision-making are complex systems that integrate many AI techniques to effectively produce cognitively congruent behavior in real time, they are computationally expensive. To support the real time simulation of an interesting demonstration scenario, we need to be able to populate the airspace with many aircraft and model multiple air traffic control sectors. To instantiate an agent on-board every aircraft and in every represented sector would dramatically reduce our potential for real-time performance. Since, at any point in the simulation, it is expected that the number of players actually engaged in a negotiation process to solve a perceived problem is likely small, we chose to implement a high-level "agent manager" entity, the purpose of which is to perform cursory conflict detection and weather analysis, and to instantiate agents to represent players that may be required to negotiate conflict resolution or weather avoidance maneuvers. When a potential conflict situation or weather hazard is detected by the agent manager, it constructs individual agent representations for the players it determines are affected by the hazard. At the very least, this includes the aircraft involved in the conflict situation. Depending on the communication structure ruleset, described previously, it may also generate an air traffic controller agent in each aircraft's current sector, or in the sector containing the perceived point of conflict. Throughout the negotiation process, agent can request that other representatives be generated to assist in solving the problem. For example, if a maneuver is proposed that would require a pilot agent to deviate significantly from its current flight plan, then that pilot might request that the agent manager construct a new agent representing its airline dispatcher. Once that agent was generated, the pilot would confer with the dispatcher agent to determine whether the proposed maneuver is acceptable.

EXPERIMENTAL APPROACH

Through the modeling effort presented here, we have developed an integrated simulation-based toolkit for the analysis of advanced distributed ATM operations. This analysis capability includes specific human behavioral representations of the key decision-makers in the future ATM environment, namely the commercial pilot, air traffic controller, and airline dispatcher. Agent-based models of their decision-making behavior within a candidate distributed decision-making framework (principled negotiation) have been developed and demonstrated. However, the integrated analysis tools have been designed and developed within a flexible architecture that allows for the extension and enhancement of the constituent modeling components within the agent representations. It is the flexibility of this architecture that provides the basis for the simulation-based toolkit for further analysis of advanced ATM operations.

One potential research question that naturally arises from the distributed decision-making approach taken in the agent development process is focused on the complexity of the predicted future ATM environment. How often will the need arise for multiple decision-makers to engage in complex negotiation to resolve perceived problems in the airspace? In standard airspace operations, most two-aircraft conflicts, for example, can be resolved through a simple even split of flight plan deviations that will generate the requisite 5 mile separation. It is expected that only in particularly complex situations with high traffic density, or significant airspace restrictions due to weather or SUA usage, will it be necessary to generate resolution maneuvers more complex than the simple even split.

To investigate this question in a simulation-based analysis, we have developed a realistic air traffic scenario based on recorded data from Jacksonville center (ZJX). This scenario is constructed from analysis of ETMS data, real-time weather track data, and SUA dynamics. The ETMS data has been pre-processed to identify specific maneuvers by individual aircraft that were implemented to avoid potential conflict or weather hazards, and original flight plan information was extrapolated from the processed information. The simulated airspace within FACET was then populated by aircraft with their original flight plans, and the recorded weather and SUA data. The scenario was run without the agents, and specific problems were identified, including traffic conflicts, weather hazards and SUA conflicts. The same scenario was then populated with our agent-based models of pilot, ATC, and airline dispatcher behavior.

First, we demonstrate that the agents are, in fact, capable of resolving these issues through distributed decision-making. Furthermore, we analyze the results of the negotiation-based process to conflict resolution, and identify the nature of the selected resolutions maneuvers. For those negotiation processes that finish in an even split of flight plan deviations across two aircraft, we conclude that in-depth negotiation was not, in fact, necessary to resolve those issues. Rather, an on-board automated system employing a sophisticated conflict resolution algorithm could have sufficed in place of multi-agent negotiation. However, for those negotiations that generate uneven flight plan deviations, for example, we conclude that automated conflict resolution algorithms may not suffice to solve these complex problems in the future ATM operational environment.

Author's Note: *We are in the process of generating a realistic air traffic scenario focused on operations in Jacksonville Center (ZJX), as described. The scenario will be based on recorded traffic data, weather data, and SUA operations over a typical period of ZJX air traffic conditions, likely spanning a period of several hours of operations. Scenario details, as well as results of post-run analysis will be presented in the final paper. In the meantime, we present results from a contrived scenario used simply to demonstrate the agents' negotiation behaviors and conflict resolution strategies.*

In the demonstration scenario, we set up a case of two aircraft flying head on towards each other and into an integrated set of weather cells of varying severity, as shown in Figure 8.

Initially, the two aircraft are not in a conflict situation. However, they are each predicted to enter unacceptable weather (the yellow and red weather cells). As a result, each pilot agent initiates a negotiation process with the ATC agent representing the first sector where a weather hazard is predicted to occur. These two separate negotiation processes are independent. Both pilots generate weather hazard avoidance maneuvers, as shown in Figure 9. However, the implementation of these weather avoidance maneuvers generates a potential traffic conflict between the two aircraft. As a result, the two pilots now negotiate a solution to this new traffic conflict, and solve it, as shown in Figure 10. Interestingly, the result of this traffic conflict resolution negotiation is a set of conflict avoidance maneuvers where one aircraft's maneuver produces 90% of the required separation, while the other produces only 10%. This is due to the added constraint imposed by the easternmost yellow weather cell. Throughout the traffic conflict resolution, the weather cells are represented as spatial constraints upon which neither aircraft may encroach. As a result, the aircraft coming from the east is not able to produce the even split in the traffic conflict resolution maneuver, and true negotiation is required to generate the final result. The events that lead to solving the weather and traffic conflicts throughout this scenario are listed in Table 1. This constitutes a simple air traffic scenario, and it is intended only for demonstration of the conflict detection and resolution capabilities of the agents through principled negotiation. We have developed more complex, yet still contrived traffic scenarios to demonstrate additional capabilities for conflict prioritization, parallel negotiation, etc.

CONCLUSIONS

We have developed an effective representation of a potential collaborative decision-making behavior for the perceived future ATM operational environment. This involved a process of knowledge elicitation (KE) with Subject Matter Experts (SMEs) to understand the issues and processes involved in current and projected ATM operations from the unique perspectives of commercial pilots, air traffic controllers and airline dispatchers. We then translated a subset of that KE information into computational models of the key cognitive processes of information processing, situation assessment and procedurally-driven decision-making. Furthermore, we integrated these cognitive processes within an overarching representation of collaborative decision-making based on principled negotiation.

Author's Note: *Additional concluding material summarizing the results of the ZJX experiment will be incorporated here.*

Under follow-on work, we recommend an effort that focuses on the further extension of the computational representations of the tasking undertaken by each of the identified key players. The selection of tasks that have been implemented under this effort was largely driven by data availability from the selected air traffic simulation environment. If additional capability was provided by the simulation, in terms of added data as well as higher-fidelity representation of aircraft dynamics, then additional cognitive modeling would become relevant and interesting. For example, the pilot and airline dispatcher have additional objectives associated with maintaining passenger comfort as well as the safety of the flight and the airspace operations. However, with the limited representation of aircraft performance, it is difficult to include these concepts effectively, and therefore, meaningfully. So, through an enhancement effort with respect to the level of fidelity provided by the simulation environment, we could set the stage for significant added capability in the cognitive modeling arena.

The focus of this effort was on generating a broad set of cognitive modeling capabilities across the key decision-makers in the advanced ATM environment. An appropriate follow-on effort should expend significant effort on the validation of these models in both a quantitative and qualitative analysis.

Finally, it has become clear that large-scale optimization of the operations within the projected collaborative ATM environment of the future will focus not only on the individual pilots, air traffic controllers and airline dispatchers, but will also include additional ATM players. Higher-level Traffic Management Unit (TMU) personnel will play a critical role. While our approach was to include a representation of these higher-level objectives within the ATC agent model, it is expected that a higher-fidelity representation of TMU objectives and behavior would have a significant effect on the overall performance of the ATM system. We have also limited our modeling capability to the enroute environment. Additional representation of the complexities of the terminal airspace is required to support appropriate end-to-end system study and design.

ACKNOWLEDGMENTS

This work was performed under NASA contract number NAS2-99051 with Ames Research Center. The authors would also like to thank Dr. Kapil Sheth and Dr. Shon Grabbe for their extensive technical support with the Future ATM Concepts Evaluation Tool (FACET).

The authors are grateful to Dr. Mica Endsley, whose expertise was invaluable throughout the knowledge engineering phases of the program. Additionally, we would like to thank our Subject Matter Experts (SMEs) for their time and patience: Lew Reszonya, Bill Jones, and Phil Bassett.

Finally, the authors thank the following Charles River Analytics staff for their technical contributions to the effort: James Omartian, Kirby Jacobs, and Raymond Pretty.

REFERENCES

Baron, S., Muralidharan, R., Lancraft, R., & Zacharias, G. L. (1980). "PROCRU: A Model for Analyzing Crew Procedures in Approach to Landing". *16th Annual Conference on Manual Control*, Moffet Field, CA. NASA Ames Research Center. pp.488-520.

Bilimoria, K., Srindhar, B., Chatterji, G., Sheth, K., & Grabbe, S. (2000). *FACET: Future ATM Concepts Evaluation Tool*.

Bilimoria, K. D. (2000). "A Geometric Optimization Approach to Aircraft Conflict Resolution". *AIAA Guidance, Navigation and Control Conference*, Denver, CO.

Denery, D. G., & Weathers, D. W. (2001). *Workshop Summary* (NASA/CP-2001-209616): NASA Ames Research Center (January).

Endsley, M., Mogford, R., Smolensky, M., & Rodgers, M. (1993). "Situation Awareness-Transition from Theory to Practice". *Proceedings of the Human Factors and Ergonomics Society 37th Annual Meeting*, Santa Monica, CA, USA. Human Factors and Ergonomics Society. pp.39-42.

Endsley, M. R., Farley, T. C., Jones, W. M., Midkiff, A. H., & Hansmann, R. J. (1998). *Situation Awareness Information Requirements for Commercial Airline Pilots* (ICAT-98-1): MIT International Center for Air Transportation

Endsley, M. R., & Rodgers, M. D. (1994). *Situation Awareness Information Requirements for En Route Air Traffic Control* (Final Report DOT/FAA/AM-94/27): Office of Aviation Medicine (December).

FAA. (1995). *Advancing Free Flight Through Human Factors* (Technical Workshop Report): Federal Aviation Administration (August 1995).

Finn, T., Labrou, Y., & Peng, Y. (1998, July). "Mobile Agents Can Benefit from Standard Efforts on Interagent Communication". *IEEE Communications Magazine*.

Guarino, S. L., Harper, K. A., White, A., Omartian, J., & Zacharias, G. (2001). *Distributed Decision-Making Models for Advanced Air/Ground Traffic Management* (Final Report R00522): Charles River Analytics Inc (August 16).

Guarino, S. L., Harper, K. A., & Zacharias, G. L. (2000). *Agent-Based Modeling & Analysis Tool for Free Flight Air Air Traffic Management: Final Report* (Final Report R97491): Charles River Analytics (9/22/2000).

Harper, K. A., Mulgund, S. M., Zacharias, G. L., & Kuchar, J. (1998). "Agent-Based Performance Assessment Tool for General Aviation Operations Under Free Flight". *AIAA Guidance, Navigation and Control Conference*, Boston, MA. Vol.1. pp.1-10.

Harper, K. A., Mulgund, S. S., Guarino, S. L., Mehta, A. V., & Zacharias, G. L. (1999). "Air Traffic Controller Agent Model for Free Flight". *Proceedings of the 1999 AIAA Guidance, Navigation and Control*, Portland, OR.

Jacolin, L., & Stengel, R. F. (1998). "Evaluation of a Cooperative Air Traffic Management Model Using Principled Negotiation Between Intelligent Agents". *AIAA Guidance, Navigation and Control Conference*, Boston, MA. Vol.1. pp.20-29.

Klein, G. A. (1989). "Recognition-Primed Decisions". In Rouse (Ed.), *Advances in Man-Machine Systems Research*, (pp. 47-92). Greenwich, CT: JAI.

Kleinman, D. L., Baron, S., & Levison, W. H. (1970). "An Optimal Control Model of Human Response, Part 1: Theory and Validation". *Automatica*, 6.

Krozel, J., Mueller, T., & Hunter, G. (1996). "Free Flight Conflict Detection and Resolution Analysis". *AIAA Guidance, Navigation and Control Conference*, San Diego, CA.

Milgram, P., van der Wijngaart, R., Veerbeek, H., Bleeker, O. F., & Fokker, B. V. (1984). "Multi-Crew Model Analytic Assessment of Landing Performance and Decision Making Demands". *Proc. 20th Annual Conference on Manual Control*, Moffett Field, CA. NASA Ames Research Center. Vol.II. pp.373-395.

NASA. (1999). *Concept Definition for Distributed Air/Ground Traffic Mangement (DAG-TM): Version 1.0* (Report): NASA (September 30).

Rasmussen, J. (1983). "Skills, Rules and Knowledge: Signals, Signs and Symbolism, and Other Distinctions in Human Performance Models". *IEEE Transactions on Systems, Man, and Cybernetics*, 12, 257-266.

Rasmussen, J. (1986). *Information Processing and Human Machine Interaction: An Approach to Cognitive Engineering*. New York: North Holland.

RTCA. (1995). *Final Report of RTCA Task Force: Free Flight Implementation* : RTCA Inc. (October 1995).

Tomlin, C., & Pappas, G. J. (1998). "Conflict Resolution for Air Traffic Management: A Study in Multi-Agent Hybrid Systems". *IEEE Transactions on Automatic Control*, 43(4), 509-521.

Wangermann, J. P., & Stengel, R. F. (1996). "Optimization and Coordination of Multi-Agent Systems Using Principled Negotiation". *AIAA Guidnce, Navigation and Control Conference*, San Diego, CA. pp.797-806.

Zacharias, G., Baron, S., Muralidharan, R., & Kastner, M. (1981). "A Supervisory Control Model of AAA Crew/System Performance and Decision Making". *17th Annual Conference on Manual Control*, UCLA, California.

Zacharias, G., Miao, A., Illgen, C., Yara, J., & Siouris, G. (1996). "SAMPLE: Situation Awareness Model for Pilot in the Loop Evaluation". *First Annual Conference on Situation Awareness in the Tactical Air Environment*, Naval Air Warfare Center, Pautexent River, MD.

Zacharias, G. L., & Baron, S. (1982). *A Proposed Crew-Centered Analysis Methodology for Fighter/Attack Missions* (4866): Bolt Beranek and Newman Inc. (February).

Figures

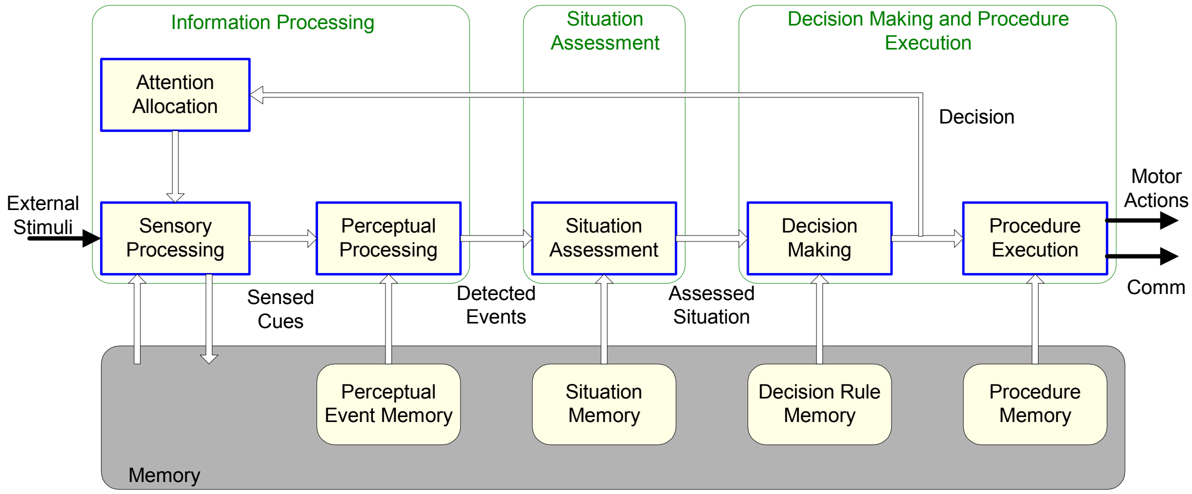

Figure 1: SAMPLE's Staged Model of Skilled Human Behavior

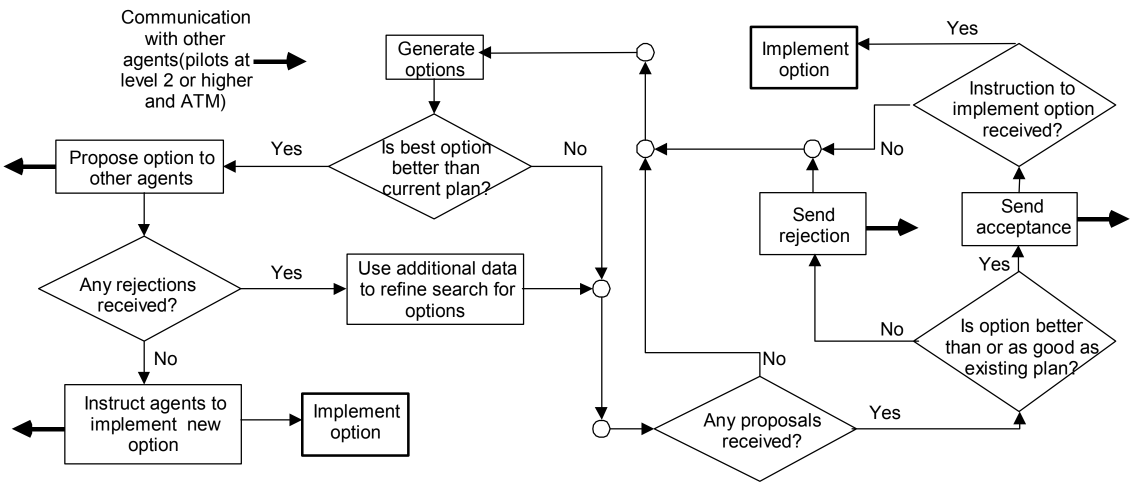

Figure 2: Principled Negotiation (Wangermann & Stengel, 1996)

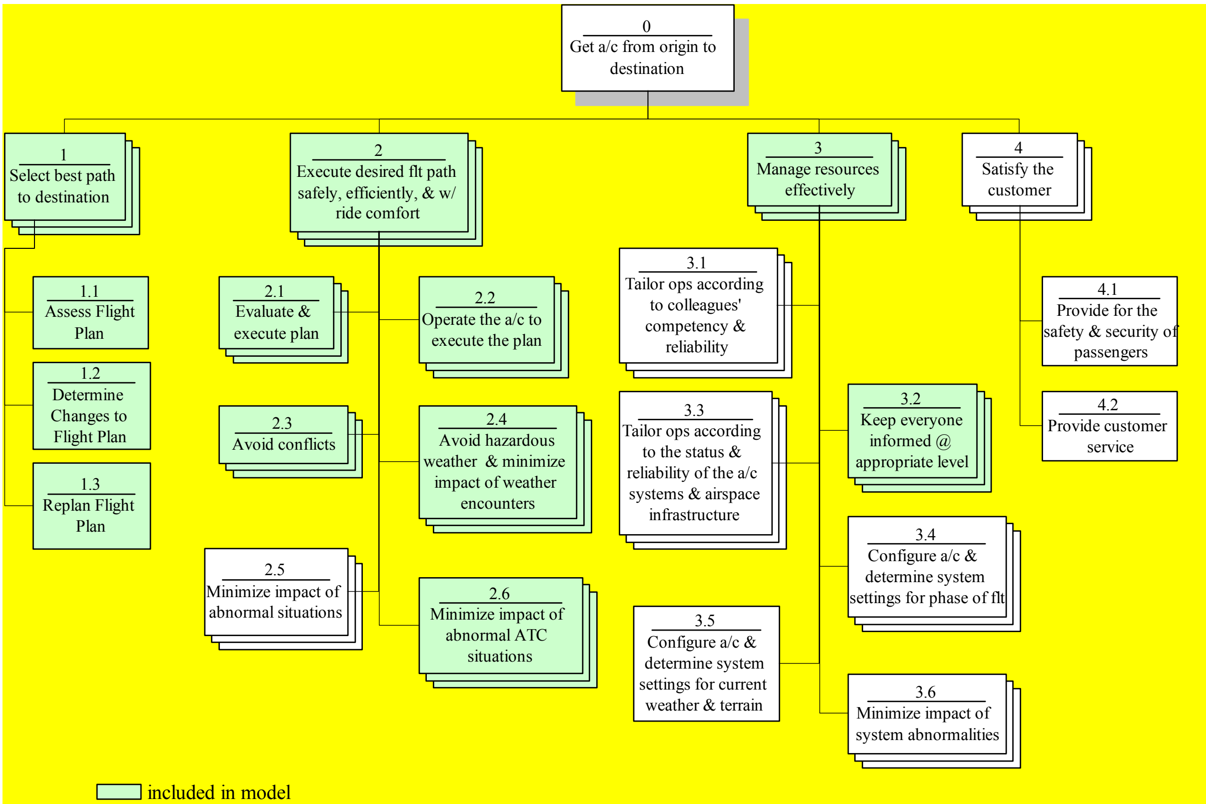

Figure 3: Top-Level Pilot Goals and Tasks

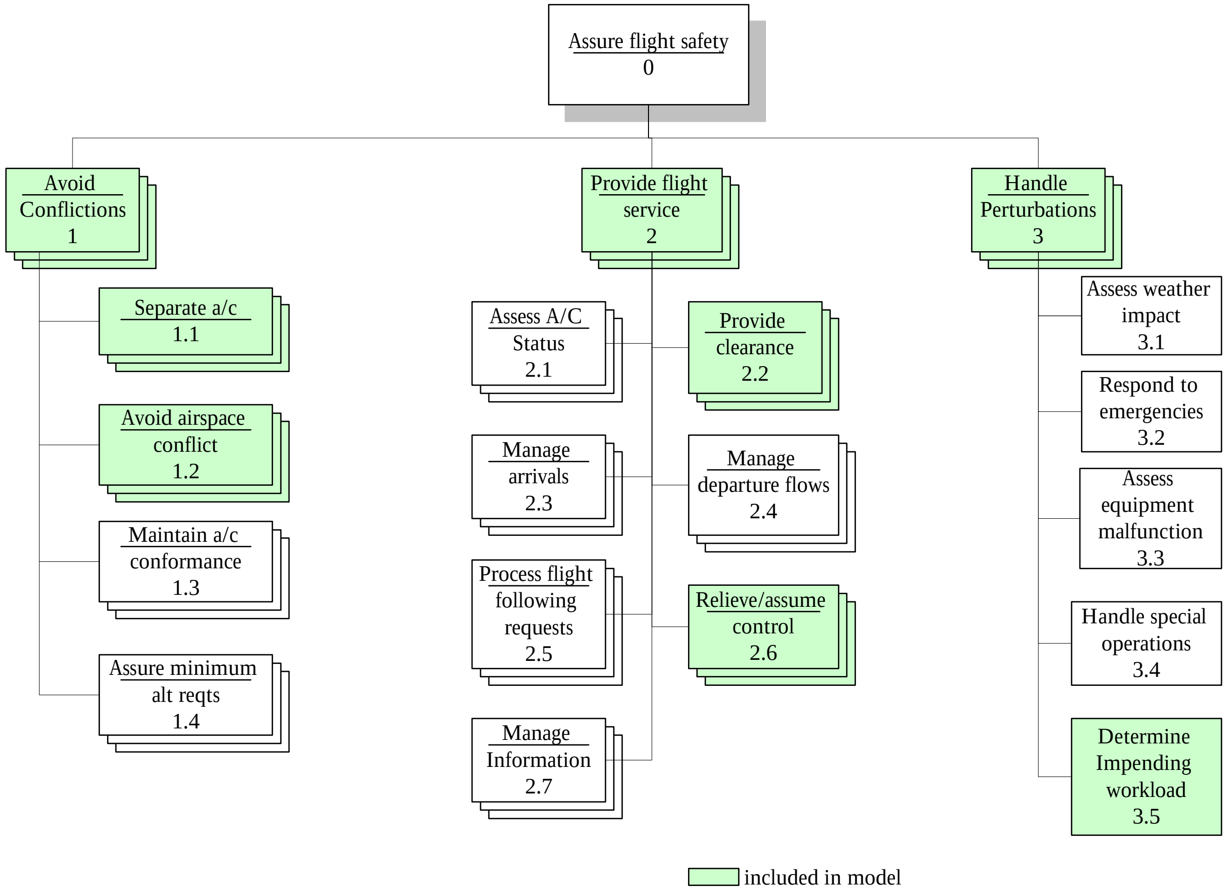

Figure 4: Top-Level ATC Goals and Tasks

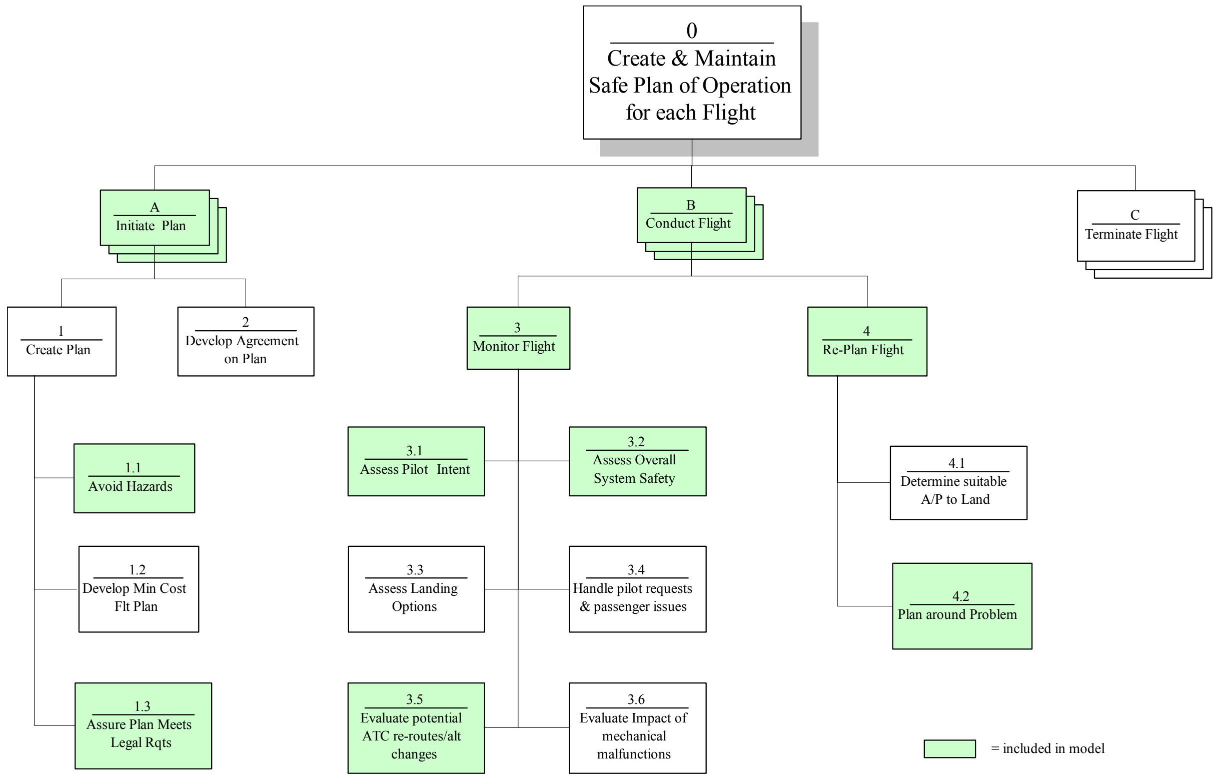

Figure 5: Top-Level Dispatcher Goals and Tasks

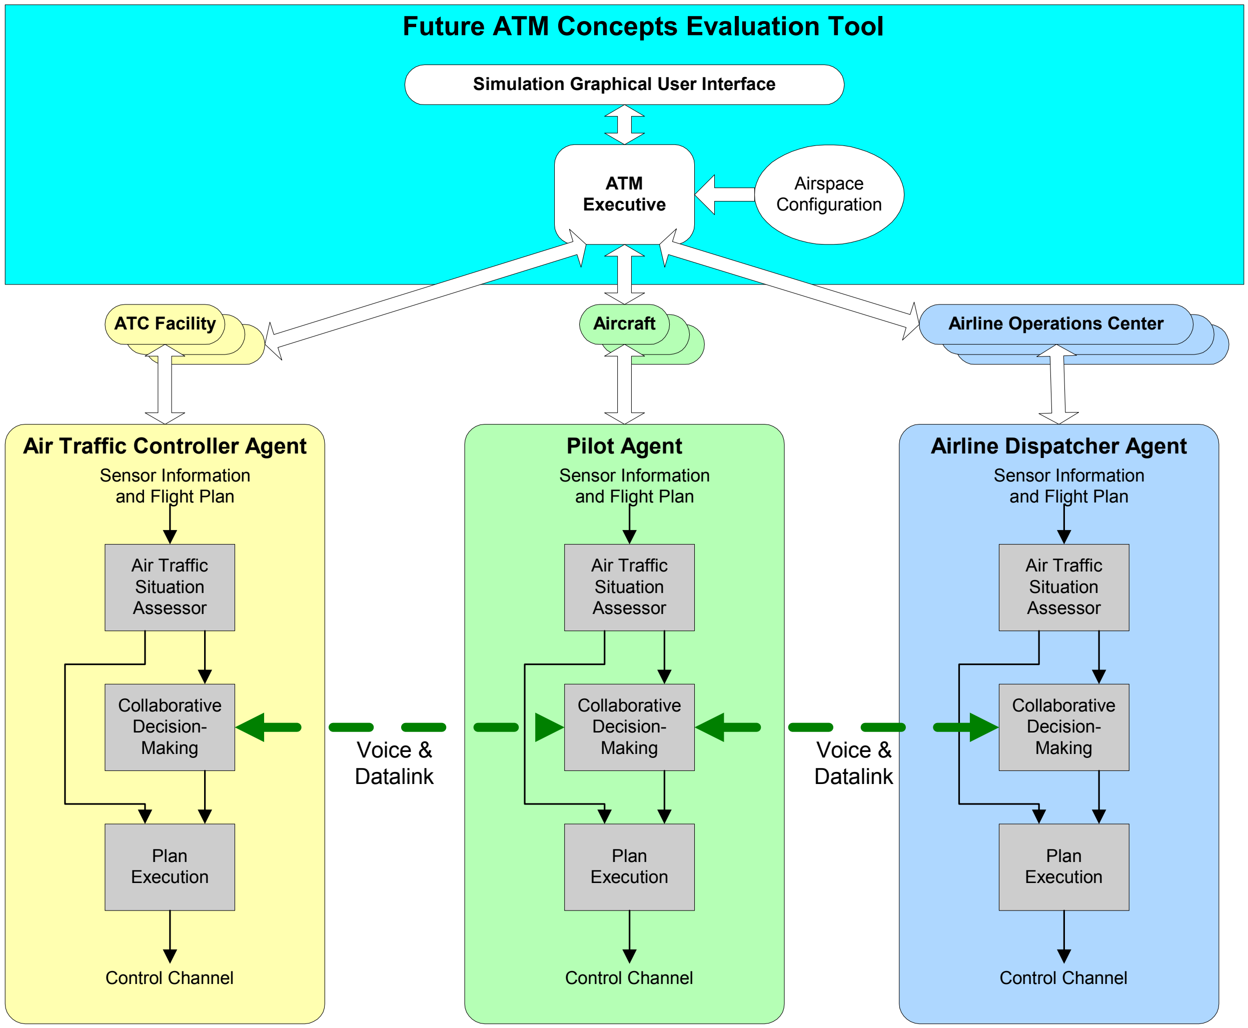

Figure 6: Overall System Architecture

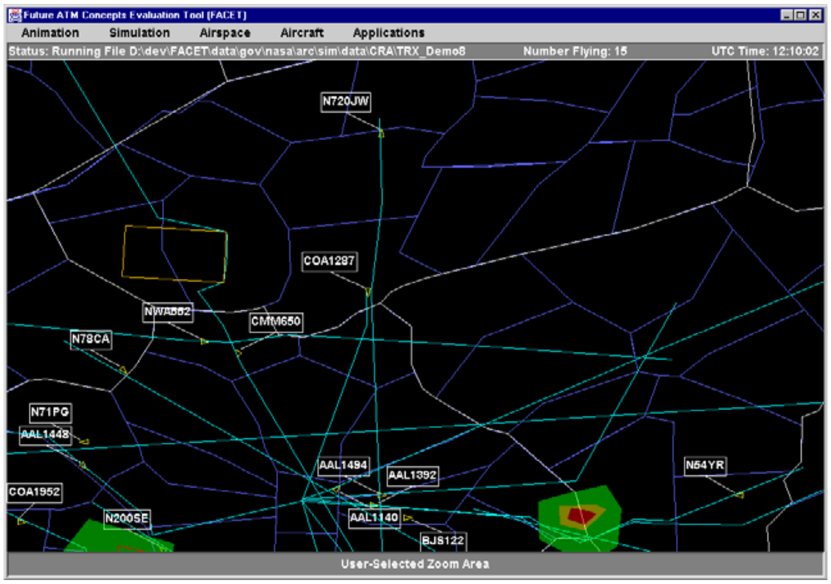

Figure 7: FACET Scenario

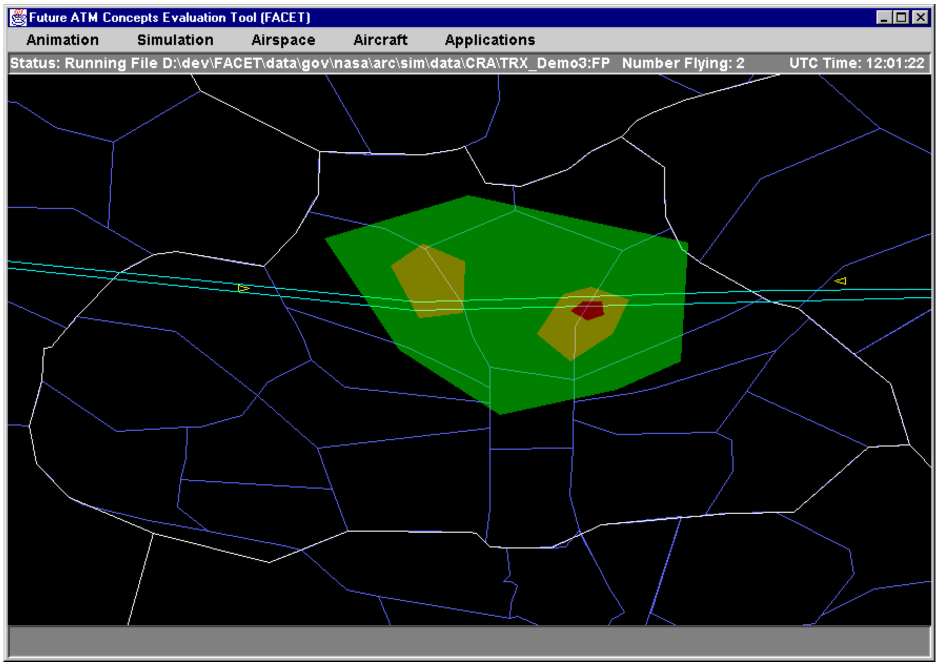

Figure 8: Initial Conditions

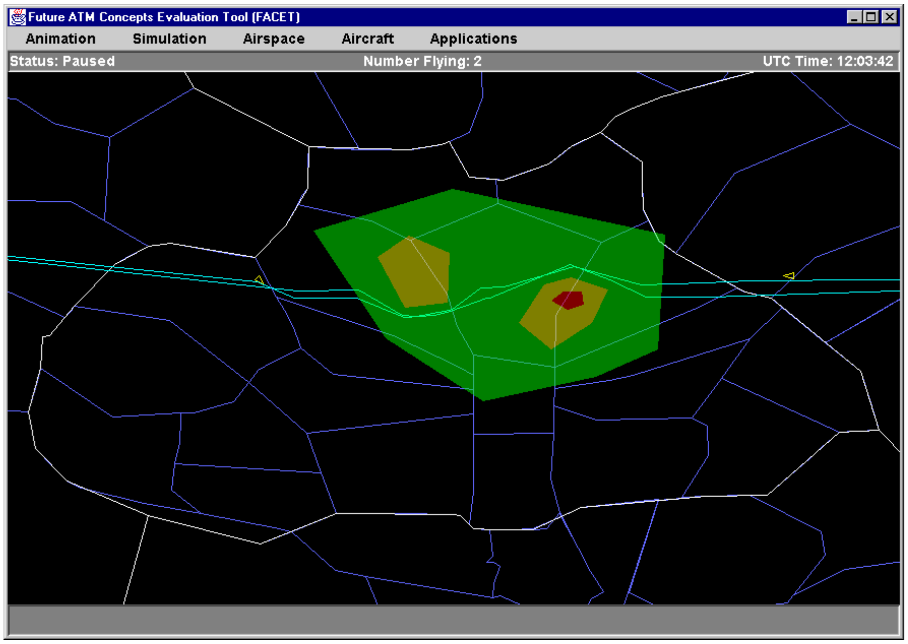

Figure 9: Weather Conflicts Solved, Causes Traffic Conflict

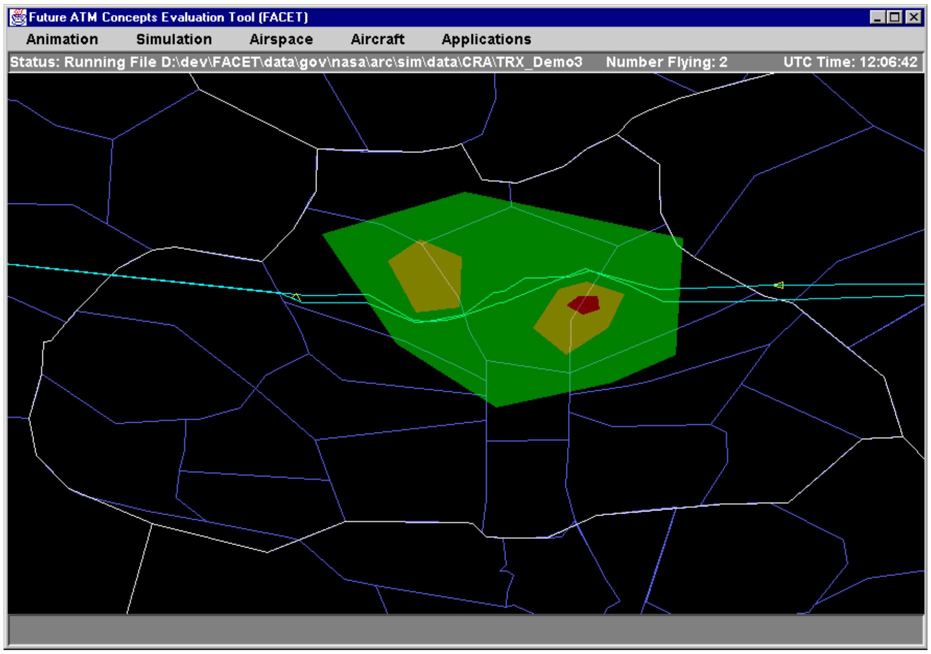

Figure 10: All Conflicts Solved

Table 1: Event Table for Solving Combined Weather/Traffic Conflict

Time	Action
1:37	NASA1 detects weather conflict and creates a negotiation to solve it.
1:41	NASA2 detects weather conflict and creates a negotiation to solve it.
---	NASA1 negotiates with ATC91 and dispatcher to reach an acceptable solution.
---	NASA2 negotiates with ATC94 and dispatcher to reach an acceptable solution.
2:59	NASA1, upon receiving approval from ATC91 and dispatcher, implements a solution to bypass the weather conflict.
3:06	NASA2, upon receiving approval from ATC94 and dispatcher, implements a solution to bypass the weather conflict.
3:35	NASA2 detects traffic conflict with NASA1 and creates a negotiation to solve it.
---	NASA2 negotiates with NASA1 and ATC95 to reach an acceptable solution.
4:44	NASA2, upon receiving approval from NASA1 and ATC95, informs NASA1 of ATC approval and implements its part of the solution.
4:45	NASA1 implements its part of the solution created by NASA2.